Backwards State-space Reduction for Planning in Dynamic Knowledge Bases

Valerio Senni
IMT Institute for Advanced Studies, Lucca, Italy
valerio.senni@imtlucca.it

Michele Stawowy
IMT Institute for Advanced Studies, Lucca, Italy
michele.stawowy@imtlucca.it

In this paper we address the problem of planning in rich domains, where knowledge representation is a key aspect for managing the complexity and size of the planning domain. We follow the approach of Description Logic (DL) based Dynamic Knowledge Bases, where a state of the world is represented concisely by a (possibly changing) ABox and a (fixed) TBox containing the axioms, and actions that allow to change the content of the ABox. The plan goal is given in terms of satisfaction of a DL query. In this paper we start from a traditional forward planning algorithm and we propose a much more efficient variant by combining backward and forward search. In particular, we propose a *Backward State-space Reduction* technique that consists in two phases: first, an *Abstract Planning Graph* $\mathscr{P}$ is created by using the *Abstract Backward Planning Algorithm* (ABP), then the abstract planning graph $\mathscr{P}$ is instantiated into a corresponding planning graph *P* by using the *Forward Plan Instantiation Algorithm* (FPI). The advantage is that in the preliminary ABP phase we produce a *symbolic* plan that is a pattern to direct the search of the concrete plan. This can be seen as a kind of informed search where the preliminary backward phase is useful to discover properties of the state-space that can be used to direct the subsequent forward phase. We evaluate the effectiveness of our ABP+FPI algorithm in the reduction of the explored planning domain by comparing it to a standard forward planning algorithm and applying both of them to a concrete business case study.

1 Introduction

In this paper we address the problem of planning in rich domains. In particular we consider the setting where a large amount of data is structured and maintained through a knowledge management system. We adopt a very well known and accepted formalism that is based on Description Logics and provides very efficient reasoning services, that is the DL-Lite framework [6].

The focus on planning comes from the industrial needs of knowledge representation and handling, as well as of work-flow management. Indeed, the large amount of data and high level of concurrency may pose a difficulty in handling resources efficiently and avoiding inconsistent behaviours. Planning can be used as a way of dynamically producing consistent work-flows on the basis of a given goal to reach. Our notion of Dynamic Knowledge Base, that changes according to actions executed by agents in the system and which we use to define our planning domain, is based on the DL-Lite framework.

Declarative work-flow management has been explored in [11] and in [12]. The first, defines a formal framework combining the constraint-based language ConDec and Computational Logic, where correctness of work-flow execution is granted by using verification techniques, run-time monitoring, and post-execution consistency analysis. The second, again based on the language ConDec, allows to specify work-flows in a declarative way by using LTL-based temporal constraints. Both of these approaches do not explicitly deal with data through an ontology as we can do using Description Logics. Another approach to the construction of correct work-flows can be found in [15], where Hierarchical Task Network

Dragan Bošnački, Stefan Edelkamp, Alberto Lluch Lafuente
& Anton Wijs (Eds.): 3rd Workshop on
GRAPH Inspection and Traversal Engineering (GRAPHITE 2014)
EPTCS 159, 2014, pp. 84–99, doi:10.4204/EPTCS.159.8

© Senni V. & Stawowy M.
This work is licensed under the Creative Commons
Attribution-No Derivative Works License.

planning is applied to a Description Logic (DL) based planning domain. This work had a different application domain with respect to ours, that is web service composition. The idea of using DLs to represent the states of the world is also used in [9], where the notion of *Knowledge and Action Base* (KAB) is introduced. Actions are applicable whenever a query has an answer in the current state of the world and has the effect of creating a new state. This last work is the one closest to ours in spirit, because it targets the business domain. However, there are several differences with our setting. First, the notion of planning domain. Second, the fact that they focus on the analysis of temporal properties of these evolving knowledge bases while we focus on the problem of state-space explosion when looking for a plan.

The contribution of this paper is the proposal of a novel planning technique that combines backwards and forwards analysis to reduce the explored state-space. This technique exploits a symbolic representation of states and reasoning techniques provided by Description Logics.

Forward search algorithms can be made very effective if an efficient search heuristics is provided. However, in lack of such heuristics they may end up exploring several paths that are not useful for reaching the goal. Backward search algorithms can take advantage of knowing the goal state to consider only actions that can lead to such a state. Our planning technique has the purpose of reducing the explored state-space by combining backward and forward planning to take advantage of both approaches. In particular, we propose a *Backward State-space Reduction* technique that consists in structuring the planning in two phases (identified jointly as ABP+FPI): first, an *Abstract Planning Graph* $\mathscr{P}$ is created by using the *Abstract Backward Planning Algorithm* (ABP), then the abstract planning graph $\mathscr{P}$ is instantiated into a corresponding planning graph P by using the *Forward Plan Instantiation Algorithm* (FPI).

The advantage of the Backward State-space Reduction technique is that, in the preliminary ABP phase, we produce (through a backward search) a *symbolic* plan that is a pattern to direct the search of the concrete plan. In particular, in the subsequent (forward) FPI construction, the abstract plan constraints the *choice of actions* to be applied and the actions have also *stronger application conditions*. This can be seen as a kind of informed search where the preliminary backward phase is useful to discover properties of the state-space that can be used to direct the subsequent forward phase. The effect of constraints induced by the abstract plan is to significantly reduce the branching of the planning by exploiting information propagated from the goal condition in a backward manner.

We have implemented the ABP+FPI algorithm and compared it to a standard Forward Planning algorithm over a case study that was designed to scale according to the values of certain parameters. The preliminary experimental results are promising both in terms of the time taken for finding the entire set of plans and in terms of the actual number of explored states.

First, we start by introducing the formalization of Dynamic Knowledge Base in Sec. 2. Then, we illustrate in Sec. 3 a Case Study which we use to test our planning technique. In Sec. 4 we describe the Backward State-space Reduction technique. Finally, in Sec. 5 we discuss a software implementation of our algorithm and some preliminary experiments.

2 Dynamic Knowledge Bases

For modelling domain resources and their relationships we adopt the Description Logic (DL) framework [2], which it is tailored to modelling a data domain by means of *concepts*, that are sets of individuals, and *roles*, that are binary relations among individuals. A DL knowledge base is made of two elements: a *TBox* T, containing axioms over concepts and roles that must hold over all individuals of the domain, and an *ABox* A, containing membership assertions of individuals to concepts and their participation into roles. Given a knowledge base $\langle T, A \rangle$, there is a number of reasoning tasks that can be

performed, among these the ones we are interested in are querying and consistency check.

Let us introduce the syntax of a fragment of DL, called DL-Lite [6]. Let N be an atomic concept name and P be an atomic role name. We can compose them by using constructors in order to define more complex concepts and roles as given in the following grammar:

$$B \;:=\; \mathsf{N} \mid \exists R \qquad C \;:=\; B \mid \neg B \qquad \text{(composed concepts)}$$
$$R \;:=\; \mathsf{P} \mid \mathsf{P}^- \qquad V \;:=\; R \mid \neg R \qquad \text{(composed roles)}$$

where $\exists$ is the projection of R over its first argument, $^-$ is the inverse relation, and $\neg$ is the complement. *TBox* axioms are constructed from (composed) concepts and roles according to the following schemes:

$$B \sqsubseteq C \quad \text{(concept inclusion)} \qquad R \sqsubseteq V \quad \text{(role inclusion)} \qquad \mathsf{funct}\, R \quad \text{(functionality)}$$

They can be translated into equivalent FOL formulas in a standard way. The *ABox* contains ground instances of concepts and roles, such as $\mathsf{N}(\mathsf{a})$ and $\mathsf{P}(\mathsf{b},\mathsf{c})$, for some individuals represented by the constants a, b, c. For the sake of simplicity, in this paper we do not allow function symbols, thus we consider only finitely many possible individuals in a knowledge base. Extensions of the DL-Lite framework allowing to reason over equalities and possibly infinite sets of individuals constructed from a finite signature are possible [1]. We plan to extend the techniques presented in this work to that more general setting.

Example 1 *We can model the hierarchy of employees within a company as follows.*

$$T \;=\; \{\mathsf{Technician} \sqsubseteq \mathsf{Employee},\ \mathsf{Manager} \sqsubseteq \mathsf{Employee},\ \mathsf{Technician} \sqsubseteq \neg \mathsf{Manager}\}$$

Where the TBox requires that technicians and managers are employees, and they are disjoint concepts (no individual can be part of both). One possible ABox is the following:

$$A \;=\; \{\mathsf{Technician}(\mathsf{e002})\}$$

where an individual identified by the constant e002 *is classified as a* Technician. *The assignment of responsibilities of managing documents is modeled by a role* assignedTo *and further axioms as follows:*

$$T' \;=\; T \cup \{\exists\, \mathsf{assignedTo} \sqsubseteq \mathsf{Document},\ \exists\, \mathsf{assignedTo}^- \sqsubseteq \mathsf{Employee},\ \mathsf{funct\ assignedTo}\}$$

Note that we set (i) the domain of the role assignedTo *to be included within* Document *(using projection), (ii) the range of the role* assignedTo *to be included within* Employee *(using projection and the inverse operator), and (iii)* assignedTo *to be functional (no document can be assigned to two employees at the same time). We will discuss an extension of this example in the case study in Sec. 3.*

We now introduce reasoning tasks over DL-Lite knowledge bases that we use in the planning algorithm: testing consistency (i.e. the existence of a model) and answering queries. A DL-Lite knowledge base is interpreted following the standard First Order Logic approach, where we fix a domain of interpretation Δ and an interpretation function $\mathscr{I}$ mapping individuals to elements of Δ, concepts to subsets of Δ, and roles to subsets of $\Delta \times \Delta$. DL-Lite axioms can be translated into equivalent FOL formulas and interpreted accordingly. An interpretation $\mathscr{I}$ is a *model* of a knowledge base $\langle T, A\rangle$ if it satisfies (the translation of) all the assertions in T and A. A knowledge base is *satisfiable* if it has a model. Finally, an *ABox* A is said to be *consistent* w.r.t. a *TBox* T if the knowledge base $\langle T, A\rangle$ is satisfiable.

Example 2 *One possible model for the knowledge base $\langle T, A\rangle$ presented in Example 1 is the one where* $\mathsf{Technician}(\mathsf{e002})$ *and* $\mathsf{Employee}(\mathsf{e002})$ *hold (obviously, considering only* $\mathsf{Technician}(\mathsf{e002})$ *would not satisfy all the axioms). An example of an inconsistent ABox is $A' = \{\mathsf{Technician}(\mathsf{e002}), \mathsf{Manager}(\mathsf{e002})\}$, which contradicts the last axiom of T.*

Queries over a DL-Lite knowledge base $\langle T, A\rangle$ are constructed considering (i) the set of the constants appearing in the assertions of A, called the *active domain* $\textsc{adom}(A)$, and (ii) the set of the predicate symbols occurring in T and A, called the *alphabet* $\textsc{alph}(T, A)$. A query q is a FOL formula of the

form $\bigvee_{i=1}^{n}(\exists y_i.\, conj_i(\vec{x}_i,\vec{y}_i))$ called *union of conjunctive queries*, where $conj_i(\vec{x}_i,\vec{y}_i)$ is a finite conjunction of atoms of the form $\mathsf{N}(z)$ and $\mathsf{P}(z,z')$, for $\mathsf{N},\mathsf{P} \in \text{ALPH}(T,A)$, and z,z' are either variables in $x_i \cup y_i$ or constants in $\text{ADOM}(A)$. The *certain answers* to a query q over $\langle T,A\rangle$ is the set $\text{ANS}(q,T,A)$ of substitutions ϑ mapping free variables of q to constants in $\text{ADOM}(A)$ and such that $q\vartheta$ holds in *every model* of $\langle T,A\rangle$, that is, $q\vartheta$ is a logical consequence of $\langle T,A\rangle$.

The most interesting feature of the DL-Lite framework is computing the set of the certain answers to a query is decidable and also efficient [1], being PTIME-complete in the size of the *ABox* and the *TBox* and AC^0 in the size of the *ABox* (this complexity is often referred to as the *data complexity*).

In some cases, it can be useful to explicitly specify a role in terms of the product of two concepts, rather than simply constraining its domain and range, as discussed in the following Example 3.

Example 3 *Let us consider the model of Example 1 and, in particular, the knowledge base $\langle T',A\rangle$ defined therein. We can further categorize documents by introducing the subclass of technical documents through the axiom* $\mathsf{TechnicalDoc} \sqsubseteq \mathsf{Document}$. *It can be useful to introduce a* canManage *role, which expresses the competences of certain employees in managing certain documents: e.g. one can be interested in modelling that* every technician can manage every technical document. *This can be used, for example, as a prerequisite for assigning documents to employees (e.g. technical documents cannot be assigned to non-technicians). However, this cannot be expressed directly as a DL-Lite axiom, since the required axiom is* $\forall\, x,y.\ (\mathsf{Technician}(x) \wedge \mathsf{TechnicalDoc}(y) \rightarrow \mathsf{canManage}(x,y))$, *which falls outside the allowed syntax. What we can model in DL-Lite is given by the following axioms:* $\exists\, \mathsf{canManage}^- \sqsubseteq \mathsf{TechnicalDoc}$ *and* $\exists\, \mathsf{canManage} \sqsubseteq \mathsf{Technician}$, *allowing only to constrain domain and range of the role of* canManage.

In general neither joins nor concept products are expressible in DL-Lite and therefore we are not allowed to consider axioms of the form $\forall\, x,y.\ (\mathsf{N}_1(x) \wedge \mathsf{N}_2(y) \rightarrow \mathsf{R}(x,y))$, as mentioned in the example. In this paper we allow *this specific form* of axioms on the basis of recent results obtained in a language called Datalog$^{\pm}$ [3]. Datalog$^{\pm}$ is indeed a family of languages, that strictly generalizes the DL-Lite family. The so-called *sticky* fragment of Datalog$^{\pm}$ [4] allows to specify concept products, as well as other more general forms of joins, and enjoys the same data complexity of DL-Lite. In general, DL-Lite axioms can be translated to Datalog$^{\pm}$ clauses. However, for the purpose of this paper, we stick to the restricted case of the DL-Lite syntax and we simply allow the use of axioms of the form $\forall\, x,y.\ (\mathsf{N}_1(x) \wedge \mathsf{N}_2(y) \rightarrow \mathsf{R}(x,y))$ (or simply $\mathsf{N}_1(x) \wedge \mathsf{N}_2(y) \rightarrow \mathsf{R}(x,y)$) in the *TBox*, keeping the good complexity results. We call these axioms *simple joins* and we assume to identify the DL-lite axioms in T by $\text{DL}(T)$ and the simple join axioms in T by $\text{SJ}(T)$. We will use this separation for the design of the planning algorithms. We call $\mathsf{R}(x,y)$ the *conclusion* of the axiom and $\mathsf{N}_1(x) \wedge \mathsf{N}_2(y)$ the *premise*. We leave for future work the extension of our planning domain reduction technique to the full Datalog$^{\pm}$ framework.

Let us now consider the *dynamical* aspect of our knowledge bases. We introduce the notion of *Dynamic Knowledge Base* (DKB), which is a transition system where states are DL-Lite knowledge bases and actions are used to update the *ABox*. In particular, we assume the *TBox* T does not change, so the ABox A is sufficient to identify the state of the system and we will feel free to refer only to A, without explicitly mentioning T. A Dynamic Knowledge Base is a tuple $\langle T,A_0,\Gamma\rangle$, where T is a TBox, A_0 is an ABox called the *initial state*, and Γ is a finite set of well-formed actions.

An *action* is of the form $\mathsf{a}[x_1,\ldots,x_n]$: $q \rightsquigarrow e$ (or simply a: $q \rightsquigarrow e$), where a is the action name, q is a query (called *action guard*), and e is an (possibly non-ground) *ABox* assertion (called *action effect*) such that $Vars(e) \subseteq Vars(q) = \{x_1,\ldots,x_n\}$. An action is *well formed* w.r.t. a knowledge base $\langle T,A\rangle$ if predicates occurring in its effects are disjoint from predicates occurring in conclusions of simple join axioms in $\text{SJ}(T)$. The well-formedness of actions is an important condition since it allows us to distinguish the assertions that are entailed by simple join axioms from the assertions that can be introduced by actions.

The informal semantics of an action is that, by using the query q, we extract *one* certain answer ϑ from the current knowledge base and we obtain a corresponding *ground ABox* assertion $e\vartheta$. The effect of an action a over a state A is, non-deterministically, a new state $A \cup \{e\vartheta\}$, for each $\vartheta \in \text{ANS}(q,T,A)$, which we indicate as $A \rightsquigarrow_{\mathsf{a},\vartheta} A \cup \{e\vartheta\}$. We call $\mathsf{a}\vartheta$ an *instantiation* of a.

Of course adding an assertion to an *ABox* could make it inconsistent. Such cases are not considered in our transition system, and actions' instantiations that lead to inconsistent *ABox*-es are ignored.

Example 4 *Following up our previous examples based on a company case study, we can model the fact that a manager can decide to assign documents to employees, e.g. for revision purposes. The following action requires to identify a manager and an employee that is able to manage a document, and under these conditions the document can be assigned to that employee:*

$\mathsf{appoint}[x,y,z]$: $\mathsf{Manager}(x) \wedge \mathsf{canManage}(y,z) \rightsquigarrow \mathsf{assignedTo}(z,y)$

Considering Example 3, if y is bound to a technical document, the only eligible employees will be technicians. The effect on the knowledge base is the addition of an assignedTo *assertion. Now, assume an ABox of the form* $A = \{\mathsf{Manager}(\mathsf{e001}), \mathsf{Technician}(\mathsf{e002}), \mathsf{Technician}(\mathsf{e003}), \mathsf{TechnicalDoc}(\mathsf{d001})\}$, *where we have two technicians. The two possible effects of the action* appoint *are*

$A \rightsquigarrow_{\mathsf{appoint},\vartheta_1} A \cup \{\mathsf{assignedTo}(\mathsf{d001},\mathsf{e002})\}$ *and* $A \rightsquigarrow_{\mathsf{appoint},\vartheta_2} A \cup \{\mathsf{assignedTo}(\mathsf{d001},\mathsf{e003})\}$,

where $\vartheta_1 = \{x \mapsto \mathsf{e001}, y \mapsto \mathsf{e002}, z \mapsto \mathsf{d001}\}$ *and* $\vartheta_2 = \{x \mapsto \mathsf{e001}, y \mapsto \mathsf{e003}, z \mapsto \mathsf{d001}\}$.
As mentioned before, actions can lead to inconsistent states:

$A \cup \{\mathsf{assignedTo}(\mathsf{d001},\mathsf{e002})\} \rightsquigarrow_{\mathsf{appoint},\vartheta_3} A \cup \{\mathsf{assignedTo}(\mathsf{d001},\mathsf{e002}), \mathsf{assignedTo}(\mathsf{d001},\mathsf{e003})\}$,

where the axiom regarding the functionality of the role assignedTo *is violated. Such transition would be discarded in the transition system.*

The notion of Knowledge and Action Base presented in [9] is in principle very similar to our DKB and has been proposed to model knowledge base dynamics. Indeed, we have been inspired by that notion, but the main difference between our definition and that one is in the way the new knowledge base is constructed as an effect of an action. In particular, in their approach, no assertion from the previous *ABox* is maintained and the new *ABox* is entirely constructed summing the effect of all possible instantiations $\mathsf{a}\vartheta$ of an action a. Furthermore, actions have also a more general form with various possible effects applied in *parallel* and *at once*. Therefore, that definition implements a semantics of the dynamics which is very different form ours. In their setting, however, our semantics can be reconstructed by explicitly adding actions that reconstruct the information that we preserve at each step. Discussions on the appropriateness of these models are longstanding and related to the frame problem as well as to other well known problems in the field of knowledge representation [13]. In this paper we do not address these problems because we focus on planning and state-space reduction aspects. We leave for future work to study how our state-space reduction technique adapt to different dynamic models.

3 Case Study

In our case study a company is interested into having a centralized control over its activities, with the aim of minimizing the effort of maintaining the consistency of the entire set of its work-flows. In the company, employees are organized in specific areas of competence and documents are assigned accordingly. In this setting, we imagine that the work of an employee is assisted by a planning software, which receives as input an high level goal and provides one or more plan to reach such goal. The planner should satisfy the company ruled represented by axioms in the *TBox*.

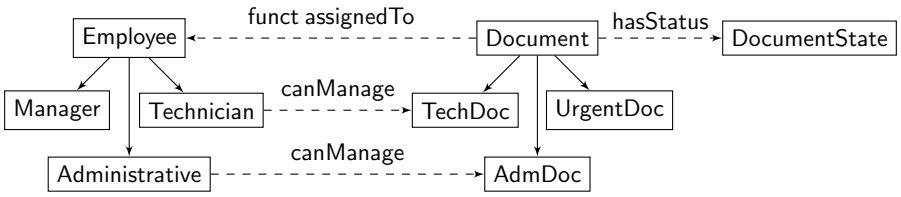

Figure 1: Graphical representation of the *TBox*

In Fig. 1 we give an informal representation of the hierarchy of employees and of documents in the company (identified by solid boxes and connected through solid lines). Roles involving these concepts of individuals are represented through dashed lines. This structure has been partially formalized by axioms discussed in examples given in Sec. 2. Now we illustrate more extensively how this structure is formalized by using DL-Lite and Simple Join axioms. The complete specification of the case study can be found in the Appendix.

The agents of the company are grouped in the concept Employee. As explained in Example 1, we build a hierarchy by introducing the concepts Manager, Administrative and Technician and using axioms of the type $\mathsf{Technician} \sqsubseteq \mathsf{Employee}$ to denote concept containment. Disjointness between concepts is specified using axioms like $\mathsf{Technician} \sqsubseteq \neg\mathsf{Manager}$. Similarly, we categorize the documents within the main concept Document, by introducing TechnicalDoc, AdministrativeDoc, and UrgentDoc, where TechnicalDoc is disjoint from AdministrativeDoc. We also create a concept DocumentState which expresses the status of a document, for example reviewed.

For the case study we consider different numbers of employees and documents by considering different sets of assertions like $\mathsf{Technician}(\mathsf{e002})$ in the *ABox*. We discuss this aspect more in details in Sec. 5, planning experiments and performance of the algorithms presented in Sec. 4 are discussed.

We define also roles in which instances can participate. The role hasStatus relates a document to a working status (like reviewed). The roles (canManage and assignedTo) describe the relationship between the employees and documents, allowing us to put some restrictions on the type of documents an employee can deal with. In particular, canManage specifies which employees can manage which documents, while assignedTo specifies which documents have been assigned to a specific employee.

The axiom $\mathsf{Technician} \sqsubseteq \exists\mathsf{canManage}.\mathsf{TechnicalDoc}$ we require that the concept Technician is a subconcept of the domain of canManage, but restricted only to the assertions in which the range are instances from the concept TechnicalDoc. Consider for example the assertion $\mathsf{canManage}(\mathsf{e004},\mathsf{d001})$: if d001 is not a technical document, then e004 cannot be a technician. We define a similar restriction over the concept Administrative, by using the axiom $\mathsf{Administrative} \sqsubseteq \exists\mathsf{canManage}.\mathsf{AdministrativeDoc}$.

On top of the simple join axiom defined in Example 3, we also introduce an axiom relative to Administrative employees and corresponding AdministrativeDoc:

$$\forall\, x,y.\ (\mathsf{Administrative}(x) \wedge \mathsf{AdministrativeDoc}(y) \rightarrow \mathsf{canManage}(x,y))$$

which defines that every pair of administrative employee and document is in the role canManage. The actions available in the set Γ are appoint (defined Example 4), review, setAdmDoc, and setTechnician. The action review allows to set a document to the reviewed state if it is in a relationship assignedTo.

$\mathsf{review}[x,y]$: $\mathsf{assignedTo}(x,y) \rightsquigarrow \mathsf{hasStatus}(x,\mathsf{reviewed})$

A manager can set a document (resp., an employee) as an administrative document (resp., a technician) through the action setAdmDoc (resp., setTechnician) given below:

$\mathsf{setAdmDoc}[x,y]$: $\mathsf{Manager}(x) \wedge \mathsf{Document}(y) \rightsquigarrow \mathsf{AdministrativeDoc}(y)$

$\mathsf{setTechnician}[x,y]$: $\mathsf{Manager}(x) \wedge \mathsf{Employee}(y) \rightsquigarrow \mathsf{Technician}(y)$

In this case study, the objective is to reach a state where an urgent document reaches the revised state,

which is expressed in terms of a *goal* as follows: $\mathsf{UrgentDoc}(x) \wedge \mathsf{hasStatus}(x, \mathsf{reviewed})$. This goal is provided to the planner for the construction of a plan that allows us to realize it.

4 Planning in Dynamic Knowledge Bases

We consider the classical reachability planning problem [8] in the domain of dynamic knowledge bases and forward and backward search algorithms. We assume *uninformed* algorithms [14], as we want them to be usable in different scenarios, where there could be no obvious heuristic function to guide the algorithm. However, our algorithms can be easily parametrized to add an appropriate search heuristic.

Given a dynamic knowledge base $\langle T, A_0, \Gamma \rangle$, formalized in Sec. 2, we define our *planning problem* as the tuple $\langle T, A_0, \Gamma, g \rangle$, where g is a query representing the *goal*. The query g is evaluated against the knowledge base $\langle T, A_i$, where A_i represents one of the states of the DKB.

Our aim is to build a *planning graph* P representing all possible plans. The planning graph P is a set of tuples $\langle A, \mathtt{a}, \vartheta \rangle$ composed of an *ABox* A, an action $\mathtt{a}$ in Γ, and a substitution ϑ used to obtain the action instance $\mathtt{a}\vartheta$. A tuple $\langle A, \mathtt{a}, \vartheta \rangle \in P$ represents the transition $A \rightsquigarrow_{\mathtt{a},\vartheta} A'$ in the system. Given a planning graph P, a plan p is a path starting from the initial state A_0 and terminating in a *goal state*, which is any state A' such that $\textsc{ans}(q, T, A') \neq \emptyset$, by using transitions (tuples $\langle A, \mathtt{a}, \vartheta \rangle$) in P. So, a plan p is a sequence $(\mathtt{a}_1\vartheta_1, \ldots, \mathtt{a}_n\vartheta_n)$ of actions instantiations. As discussed in Sec. 2, actions could lead to inconsistent states (i.e. *ABox*-es w.r.t. the *TBox* T), thus we need to take into account this aspect in the construction of the planning algorithms.

As said before, by default we want to find all possible plans that are an answer to the planning problem. Such choice is justified by the contest in which we want to operate, namely business rich domains, where the decision about which plan to follow could be demanded to other components than the planner itself. Another reason, more related to our implementation, is that our actions do not have any cost function attached to them, thus making it risky to choose one plan other another basing the decision on classic metrics like the number of actions. Anyway, the planning algorithms shown below can be easily parametrized to find all or just one plan.

4.1 Forward Planning

Forward Planning starts from an initial state and expands it by applying all executable actions in that state. New states obtained in this way are added to the states that must be explored, except for those that are inconsistent or are goal states. Actions are executable in a state if their guard has at least one answer in that state. For a given state and action, we have as many outgoing edges as the number of answers. The *Forward Planning Algorithm* (FP) is presented in a pseudo-code fashion in Algorithm 1.

The algorithm FP takes as input a dynamic knowledge base $\langle T, A_0, \Gamma \rangle$ and a goal g, representing the property of state to be reached, and returns a planning graph P as output. FA operates on two sets of states: R (the *remaining* states), containing states that have to be expanded, and V (the *visited* states), containing states that have already been expanded. R is initialized to $\{A_0\}$, which is the initial state. Recall that states are represented by mentioning explicitly only the *ABox*, since the *TBox* is constant.

The main loop of the algorithm takes one state from R, adds it to V, checks if A is consistent, and if so, proceeds to expand that state. If A is inconsistent it is simply eliminated, but kept in the set of visited states to avoid further consistency tests over it. We assume to have a function `Consistent`, with boolean return value, that allows us to test the consistency of an *ABox* w.r.t. a *TBox*. On elimination of A, also the edges in P that lead to A must be eliminated, which is done by using the auxiliary function `EdgesTo`,

defined as follows:

$$\texttt{EdgesTo}(P,\Gamma,A) \;=\; \{\langle A',\mathsf{a},\vartheta\rangle \mid \langle A',\mathsf{a},\vartheta\rangle \in P \;\wedge\; (\mathsf{a}\colon \mathsf{q} \rightsquigarrow \mathsf{e}) \in \Gamma \;\wedge\; A = A' \cup \{e\vartheta\}\}$$

This function removes from P the edges that can cause the generation of the inconsistent state A. On the contrary, if A is a consistent state and also a goal state ($\textsc{ans}(q,T,A') \neq \emptyset$) it is not expanded further. Otherwise, if A is not a goal state, it is expanded by finding *all* of the outgoing edges as well as the corresponding arrival states, which is done using the auxiliary function `Next`, defined as follows:

$$\texttt{Next(T,A,}\Gamma\texttt{)} \;=\; \{\langle \mathsf{a},\vartheta,A'\rangle \mid (\mathsf{a}\colon q \rightsquigarrow e) \in \Gamma \wedge \vartheta \in \textsc{ans}(q,T,A) \wedge A \rightsquigarrow_{\mathsf{a},\vartheta} A'\}$$

This function considers every action $\mathsf{a}\colon q \rightsquigarrow e$ and every action instance obtained by considering an answer ϑ in the set $\textsc{ans}(q,T,A)$, where q is the action guard. For every such action a and answer ϑ, it expands A into the new state A' by the action $A \rightsquigarrow_{\mathsf{a},\vartheta} A'$. The new states obtained through the function `Next` are added to R, except for those that have already been visited (i.e. the states in V). The consistency of these new states will be tested when they will be selected for expansion from the set R. Finally, the *Planning Graph* P is updated by adding the new edges found through the function `Next`.

Algorithm 1: Forward Planning Algorithm

input : A dynamic knowledge base $\langle T,A_0,\Gamma\rangle$ and a goal g, with A_0 consistent w.r.t. T
output: A planning graph P

```
begin
    P := ∅                                                    (planning graph)
    R := {A_0}                                                (remaining states)
    V := ∅                                                    (visited states)
    while R ≠ ∅ do
        A ∈ R
        R := R \ {A}
        V := V ∪ {A}
        if ¬Consistent(A,T) then
            P := P \ EdgesTo(P,Γ,A)                           (remove edges reaching A)
        else if ANS(g,T,A) ≠ ∅ then
            skip
        else
            R := R ∪ ( {A' | ⟨a,ϑ,A'⟩ ∈ Next(T,A,Γ)} / V )    (add new states to R)
            P := P ∪ {⟨A,a,ϑ⟩ | ⟨a,ϑ,A'⟩ ∈ Next(T,A,Γ)}        (update plan)
        end
    end
end
```

Note that the FP algorithms finds *all* possible plans, because the retrieval of a goal state does not interrupt the execution of the loop, which terminates only when all states have been visited. The variant of the algorithm where only *one* plan is found can be easily obtained by replacing the **skip** command, after the test on being a goal state, with a **break** command, which interrupts the execution of the loop thus skipping the visit of all of the remaining states.

In both cases the FP algorithm always terminates because the set of all possible states (*ABox*-es) is finite, since (i) the set of all possible individuals is finite and consists in the $\textsc{adom}(A_0)$ plus the individuals occurring in the set of actions Γ, (ii) the set of all possible concepts and roles is finite as well, and (iii) actions cannot create new individuals, nor new predicate symbols.

As usual in graph search algorithms, different search strategies (such as breadth-first or depth-first)

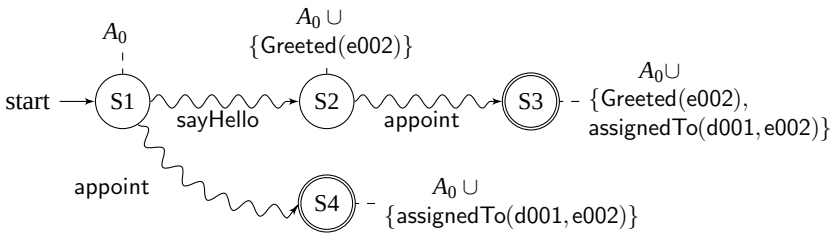

Figure 2: Forward Algorithm example, Planning Graph

as well as heuristics can be added by imposing a data structure on the set R and a specific extraction strategy. Thus the algorithm FP can easily be adapted for an *informed* search.

The forward algorithm has drawbacks due to the fact that, being uninformed, it explores all feasible plans. In particular, it produces also *redundant plans*, that are plans containing proper sub-plans. We define a plan $(\mathsf{a}_1\vartheta_1)\ldots(\mathsf{a}_n\vartheta_n)$ to be *redundant* if there exists a proper sub-sequence of $(\mathsf{a}_1\vartheta_1)\ldots(\mathsf{a}_n\vartheta_n)$ which is itself a plan. An example of redundant plan is discussed below.

Example 5 *Let us consider the Dynamic Knowledge Base $\langle T_0, A_0, \Gamma_0\rangle$ where:*

$$T_0 = T'$$
$$A_0 = A \cup \{\mathsf{Manager}(\mathsf{e001}), \mathsf{TechnicalDoc}(\mathsf{d001}), \mathsf{canManage}(\mathsf{e002}, \mathsf{d001})\}$$
$$\Gamma_0 = \Gamma \cup \{\mathsf{sayHello}[x,y]\text{: } \mathsf{Manager}(x) \wedge \mathsf{Technician}(y) \rightsquigarrow \mathsf{Greeted}(y)\}$$

The knowledge base $\langle T', A\rangle$ is taken from Example 1, and Γ contains the action appoint *defined in Example 4). We are given the goal:* $g \leftarrow \mathsf{assignedTo}(\mathsf{d001}, \mathsf{e002})$

The FA would produce the following steps:
(1) select A_0 as the first state to expand (it is consistent, but it is not a goal state since $\text{ANS}(g, T_0, A_0) = \emptyset$),
(2) compute $\langle \mathsf{appoint}, \{x \mapsto \mathsf{e001}, y \mapsto \mathsf{e002}, z \mapsto \mathsf{d001}\}, A_1\rangle$, where $A_1 = A_0 \cup \{\mathsf{assignedTo}(\mathsf{d001}, \mathsf{e002})\}$, and $\langle \mathsf{sayHello}, \{x \mapsto \mathsf{e001}, y \mapsto \mathsf{e002}\}, A_2\rangle$, where $A_2 = A_0 \cup \{\mathsf{Greeted}(\mathsf{e002})\}$, using the function `Next`,
(3) update the set R and the planning graph P accordingly,
(4) select from R the state A_1, which is consistent and is also a goal state, thus is skipped,
(5) expand the state A_2, which is consistent, but $\text{ANS}(g, T_0, A_2) = \emptyset$, so it is not a goal state),
(6) compute $\langle \mathsf{appoint}, \{x \mapsto \mathsf{e001}, y \mapsto \mathsf{e002}, z \mapsto \mathsf{d001}\}, A_3\}\rangle$, where $A_3 = A_2 \cup \{\mathsf{assignedTo}(\mathsf{d001}, \mathsf{e002})\}$, using the function `Next`
(7) update the set R and the planning graph P accordingly,
(8) select from R the state A_3, which is consistent and is also a goal state, thus it is skipped,
(9) terminate because R is empty.
We show a graphical representation of the planning graph in Fig. 2 (for simplicity, we do not mention the substitutions). The set of plans represented by P is very simple: $\{(\mathsf{sayHello}, \mathsf{appoint}), (\mathsf{appoint})\}$. In particular, the plan $(\mathsf{sayHello}, \mathsf{appoint})$ is redundant, since it contains the proper sub-plan $(\mathsf{appoint})$.

On large examples, the impact of redundant plans is a much larger planning graph. The backward planning technique we are going to present in the rest of this section has the purpose of reducing the number of redundant plans found and, thus, the size of the planning graph. The application of the Forward Planning Algorithm to the Case Study (Sec. 3) is discussed in Sec. 5.

4.2 Backward State-space Reduction and Planning

Backward Planning [7] is based on the idea of starting from a goal state and, analysing actions that can lead to such state, build the set of predecessor states until either the initial state is found or the state-space is explored entirely, thus being *goal-driven*. In lack of specific information on the state-space structure,

backward algorithms can have advantages over forward algorithms since they explore a smaller portion of the state-space by considering only actions that can lead to the goal satisfaction. On the contrary, in the forward approach and under uninformed search, all the executable actions are taken into account and more states are generated. Another very common and widely used technique to make the planning problem easier to solve is to create an *abstraction* of it [14]: the original planning domain is shrank by omitting details, thus reducing the size of it (by removing superfluous states) and making it more manageable, but without loosing the capability of finding plans.

This is also what we do in the proposed *Backward State-space Reduction* technique, where we combine abstraction and backward state-search. The idea is to structure the planning in two phases (identified jointly as ABP+FPI): first, an *Abstract Planning Graph* $\mathscr{P}$ is created by using the *Abstract Backward Planning Algorithm* (ABP), then the abstract planning graph $\mathscr{P}$ is instantiated into a corresponding planning graph P by using the *Forward Plan Instantiation Algorithm* (FPI), which is essentially a variant of FP. The abstraction we apply in ABP is that the algorithm manipulates states represented by *queries* rather than *ABox*-es. A query q (which we call an *abstract state*) represents a set *concrete*(q,T) of *ABox*-es such that $A \in$ *concrete*(q,T) iff (i) A is consistent w.r.t. T and (ii) $\textsc{ans}(q,T,A) \neq \emptyset$. In this sense, the algorithm is *symbolic* and manipulates sets of (possible) states rather than single states. Having the abstract planning domain, the ABP produces (through a backward search) an *abstract planning graph* $\mathscr{P}$, which is a set of tuples $\langle \sigma, \mathsf{a} \rangle$ composed of an abstract state σ and an action a: $q \rightsquigarrow e$ in Γ. The tuple $\langle \sigma, \mathsf{a} \rangle \in P$ represents *the set of all transitions* $A \rightsquigarrow_{\mathsf{a},\vartheta} A'$, such that (i) $A \in$ *concrete*(σ,T), and (ii) $\vartheta \in \textsc{ans}(q \wedge \sigma, T, A)$. In other words, we interpret $\langle \sigma, \mathsf{a} \rangle$ as a constraint over the guard q of action a, which is refined using the abstract state (i.e. query) σ.

The abstract planning graph $\mathscr{P}$ is then used within the FPI algorithm, which is essentially a variant of FP, as a pattern to direct the search of the concrete plan P. In particular, we force the choice of actions to be applied and the actions have also stronger guard, as we shall discuss in the following. The advantage of enforcing constraints over actions, w.r.t. the FP algorithm, is to significantly reduce the branching of the planning by exploiting information propagated from the goal condition in a backward manner. This can be seen as a kind of informed search where the preliminary abstract backward phase is useful to discover properties of the state-space that can be used to direct the second phase.

We now describe the two algorithms in detail. The Abstract Backward Planning Algorithm is given in pseudo-code in Algorithm 2. It keeps track of remaining abstract states to be explored in the set R and of visited abstract states in the set V. For every abstract state σ in R, the algorithm tests whether it includes the initial *ABox* A_0, if so the abstract state is not further expanded. Otherwise, the abstract state σ is *resolved* to a new set of abstract states by applying as much as possible the simple join axioms contained in the *TBox* T (they are denoted by $\textsc{sj}(T)$). This step is performed using the auxiliary function:

$$\texttt{FullyResolve}(\sigma, \textsc{sj}(T)) = \{\, \sigma' \mid \sigma' \text{ obtained by applying } \texttt{Resolve} \text{ to } \sigma \text{ as much as possible using } \textsc{sj}(T) \text{ axioms} \,\}$$

which relies on the `Resolve` function, that applies a form of SLD-resolution in the spirit of Logic Programming [10]. In particular, a new abstract state σ' is obtained from the abstract state σ using an axiom $\mathsf{N}_1(x) \wedge \mathsf{N}_2(y) \rightarrow \mathsf{R}(x,y)$ in $\textsc{sj}(T)$ and replacing a conjunct of the form $(a_1 \wedge \ldots \wedge \mathsf{R}(x,y) \wedge \ldots \wedge a_n)\vartheta$ in σ, for some substitution ϑ, with the corresponding conjunct $(a_1 \wedge \ldots \wedge \mathsf{N}_1(x) \wedge \mathsf{N}_2(y) \wedge \ldots \wedge a_n)\vartheta$. In principle, we may have many axioms with the same conclusion and this is the reason why the function `FullyResolve` returns a *set* of possible resolved abstract states. The importance of this step is due to the fact that atoms in the conclusions of simple axioms have disjoint predicates from atoms in the effect of the actions: by resolving σ w.r.t. simple join axioms we enable the (backward) application of more actions than those applicable directly in σ. On the contrary, in forward planning we have a concrete

Algorithm 2: Abstract Backward Planning Algorithm

input : A dynamic knowledge base $\langle T, A_0, \Gamma \rangle$ and a goal g
output: An abstract planning graph $\mathscr{P}$

begin
- $\mathscr{P} := \emptyset$ (abstract planning graph)
- $R := \{g\}$ (remaining abstract states)
- $V := \emptyset$ (visited abstract states)
- **while** $R \neq \emptyset$ **do**
  - $\sigma \in$ R
  - $R := R \setminus \{\sigma\}$
  - $V := V \cup \{\sigma\}$
  - **if** $\text{ANS}(\sigma, T, A_0) \neq \emptyset$ **then** (test if σ is satisfied in the initial state)
    - **skip**
  - **for** σ' **in** `FullyResolve`$(\sigma, \text{SJ}(T))$ **do** (apply simple join axioms)
    - $R := R \cup (\{\sigma'' \mid \langle \sigma'', \mathsf{a} \rangle \in$ `PrevA`$(\sigma', \Gamma)\} \setminus V)$ (update remaining states)
    - $\mathscr{P} := \mathscr{P} \cup$ `PrevA`(σ', Γ) (update plan)
  - **end**
- **end**

end

ABox, the set of applicable actions depends on that *ABox*, and the simple join axioms do not compare explicitly in the algorithm as they are used, implicitly and together with the other axioms in the *TBox*, to test the satisfiability of the action guard.

For every resolved abstract state σ' the algorithm computes the set of previous abstract states by considering all possible actions in Γ that have an effect that is (unifiable with) an atom a in σ' and replacing a with the corresponding (unified) action guard. This is performed by the auxiliary function:

$$\texttt{PrevA}(\sigma, \Gamma) = \{\langle \sigma', \mathsf{a} \rangle \mid (\mathsf{a}\colon q \rightsquigarrow e) \in \Gamma \wedge \sigma' \in \texttt{ActPrevA}(\sigma, (\mathsf{a}\colon q \rightsquigarrow e))\}$$

$$\texttt{ActPrevA}(\sigma, (\mathsf{a}\colon q \rightsquigarrow e)) = \{\sigma' \mid \sigma' \text{ obtained by applying } \texttt{Resolve} \text{ to } \sigma \text{ using } \mathsf{a}\colon q \rightsquigarrow e\}$$

which, again, relies on the `Resolve` function that computes the new abstract state σ' by using the action $\mathsf{a}\colon q \rightsquigarrow e$ and replacing a conjunct of the form $(a_1 \wedge \ldots \wedge e \wedge \ldots \wedge a_n)\vartheta$ in σ, for some substitution ϑ, with the corresponding conjunct $(a_1 \wedge \ldots \wedge q \wedge \ldots \wedge a_n)\vartheta$.

Termination of the ABP algorithm follows from the fact that we can have finitely many abstract states (i.e. queries), for similar reasons to those discussed for termination of the FP algorithm.

Example 6 *Let us consider the DKB and the goal defined in Example 5.*
The ABP would produce the following steps:
(1) the goal state g is not satisfied in the initial KB ($\text{ANS}(goal, T_0, A_0) = \emptyset$*) and it is kept for expansion,*
(2) the set $\text{SJ}(T)$ *of simple join axioms is empty, so the function* `FullyResolve` *returns the state g itself,*
(3) the function `PrevA` *computes the pair* $\langle S_2, \mathsf{appoint} \rangle$*, where* $S_2 = \mathsf{Manager}(x) \wedge \mathsf{Technician}(\mathsf{e002})$*,*
(4) the set R and the abstract planning graph $\mathscr{P}$ *are updated accordingly,*
(5) S_2 *is selected from R and, since it is satisfied in the initial KB (*$\text{ANS}(S_2, T_0, A_0) \neq \emptyset$*), it is not expanded,*
(6) the set R is empty and the algorithm terminates. We show the resulting abstract graph in Fig. 3.

The Forward Plan Instantiation Algorithm is given in pseudo-code in Algorithm 3. This algorithm is very similar to the FP algorithm: it takes as input a dynamic knowledge base, a goal, and also an abstract planning graph. The algorithm differs from FP in the fact that actions form Γ are executed under the constraints present in $\mathscr{P}$. The plan construction starts from the initial state A_0 and each new state A in R

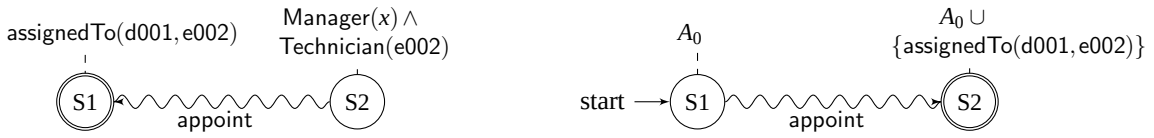

Figure 3: ABP+FPI example: Abstract Planning Graph (left) and Instantiated Planning Graph (right)

Algorithm 3: Forward Plan Instantiation Algorithm

input : A dynamic knowledge base $\langle T, A_0, \Gamma \rangle$, an abstract planning graph $\mathscr{P}$, and a goal g
output: A planning graph P

```
begin
    P := ∅                                                          (planning graph)
    R := {A0}                                                       (remaining states)
    V := ∅                                                          (visited states)
    while R ≠ ∅ do
        A ∈ R
        R := R \ {A}
        V := V ∪ {A}
        if ¬Consistent(A,T) then
            P := P \ EdgesTo(P, A)                                  (remove edges reaching A)
        else if ANS(g, T, A) ≠ ∅ then
            skip
        else
            R := R ∪ ( {A' | ⟨a, ϑ, A'⟩ ∈ NextA(T,A,𝒫,Γ)} / V )     (add new states to R)
            P := P ∪ { ⟨A, a, ϑ⟩ | ⟨a, ϑ, A'⟩ ∈ NextA(T,A,𝒫,Γ) }    (update plan)
        end
    end
end
```

is expanded by considering pairs $\langle \sigma, \mathsf{a} \rangle \in \mathscr{P}$ such that the abstract state σ includes A and the action a is executed under the extra precondition σ. This is performed by the following auxiliary function:

$$\mathtt{NextA}(T, A, \mathscr{P}, \Gamma) \;=\; \{\langle \mathsf{a}, \vartheta, A' \rangle \mid \langle \sigma, \mathsf{a} \rangle \in \mathscr{P} \wedge\; A \in \mathit{concrete}(\sigma, T) \wedge (\mathsf{a}\colon q \rightsquigarrow e) \in \Gamma \;\wedge\; \vartheta \in \textsc{ans}(q \wedge \sigma, T, A) \wedge A \rightsquigarrow_{\mathsf{a},\vartheta} A' \}$$

which consider effects $A \rightsquigarrow_{\mathsf{a},\vartheta} A'$ over the current state A with an instantiation of the action a that is computed in the set $\textsc{ans}(q \wedge \sigma, T, A)$ of answers restricted with the extra precondition σ.

The termination of the FPI algorithm is granted by the same observations make for the FP algorithm. Soundness is ensured by the fact that the algorithm find a (not-necessarily proper) subset of the plans found by the FP algorithm. Concerning completeness, we are currently working on the notion of non-redundant plan and on a relative completeness result, stating that the ABP+FPI algorithm is complete w.r.t. the set of non-redundant plans computed by the FP algorithm.
Let us now show the FPI algorithm at work.

Example 7 *Consider the DKB and the goal defined in Example 5 and the Abstract Planning Graph obtained in Example 6. The FPI produces the following steps:*
(1) A_0 *is taken as the first state to be expanded, since it is consistent but it is not a goal state,*
(2) using the function `NextA`*, find the tuple* $\langle \mathsf{appoint}, \{x \mapsto \mathsf{e001}, y \mapsto \mathsf{e002}, z \mapsto \mathsf{d001}\}, A_1 \rangle$*, where* $A_1 = A_0 \cup \{\mathsf{assignedTo}(\mathsf{d001}, \mathsf{e002})\}$
(3) the set R and the planning graph $\mathscr{P}$ *are updated accordingly,*
(4) the state A_1 *is selected from R, it is consistent w.r.t.* T *and it is a goal state (*$\textsc{ans}(q, T_0, A_1)$*), (5) the*

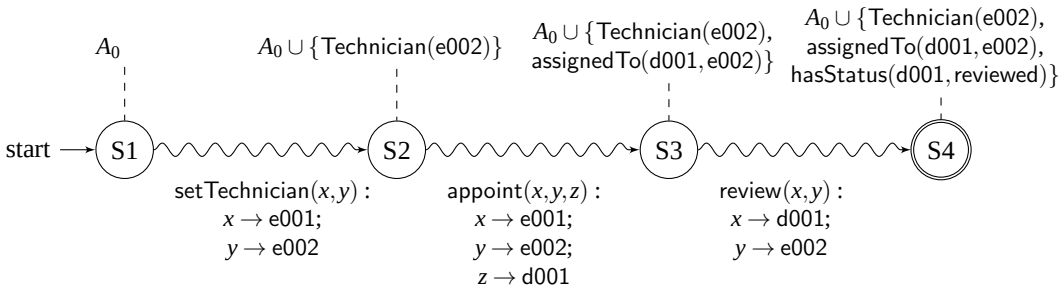

Figure 4: Planning Graph

set R is empty, so the algorithm terminates. We show the resulting graph in Fig. 3(for simplicity, we omit the substitution function ϑ).

As we can see from Example 7, the ABP+FPI produces a smaller final graph if compared to Figure 2.

5 Experiments

We have implemented the algorithms presented in Sec. 4 and in this section we report on the implementation as well as on empirical results obtained by applying the algorithms to the case study of Sec. 3. The implementation, made with Python, provides the FP algorithm or the ABP+FPI algorithm, specifying for each of them the strategy preferred (depth- or breadth-first, only the first solution or all of them).

Since the DKB is based on DL-Lite fragment, we need a reasoner to check consistency and querying knowledge bases. The reasoner Mastro [1] supports DL-Lite but it is still in closed beta and works mainly as a stand alone system. Furthermore, Mastro does not support reasoning over Simple Join axioms. Since DL-Lite is a subset of the *Web Ontology Language* (OWL) [5], we resort to the the reasoner Pellet[2], a popular and freely available OWL2 reasoner, that has all the features interested in. In particular, Pellet supports SJ axioms encoded as SWRL Rules [16]. Provided the reasoner satisfies our requirements, its choice is not crucial for our planning algorithm, since it is parametric w.r.t. the chosen reasoner.

To test the two algorithms, we created various *ABox*-es differing only for the number of instances and we varied the number of instances participating in the classes Manager, Employee and TechnicalDoc. This affects the size of the planning search space and it is useful to assess how the algorithms scale.

In Fig. 4 we present the Planning Graph obtained with FP (the instantiated planning graph obtained with ABP+FPI is identical), considering 1 manager, 1 employee and 1 technical document. For this small problem instance the two algorithms (FP and ABP+FPI) take the same amount of time. This can be explained by looking at the column *Inc* in Table 1, counting the number of inconsistent states that each the algorithm finds. Even for such a simple example, FP finds 13 inconsistent states, while FPI, thanks to the constrains provided by the Abstract Planning Graph, finds only 3 inconsistent states.

In Fig. 5 we show the *Abstract Planning Graph* obtained by the ABP (gray states are initial states). The Simple Join axioms applications are explicitly shown to make it easier to understand the algorithm behaviour. The graph is *constant* for all the *ABox*-es we created, because no matter what is the number of instances in the *ABox*, the abstract graph is always the same. Looking at Table 1, we can see that the number of states (in the Planning Graph and inconsistent ones) is greatly reduced with respect to FP.

The abstract graph may show plans that cannot be found in the instantiated planning graph. As an example of this, consider the following *ABox*:

[1] http://www.dis.uniroma1.it/~mastro/

[2] http://clarkparsia.com/pellet/

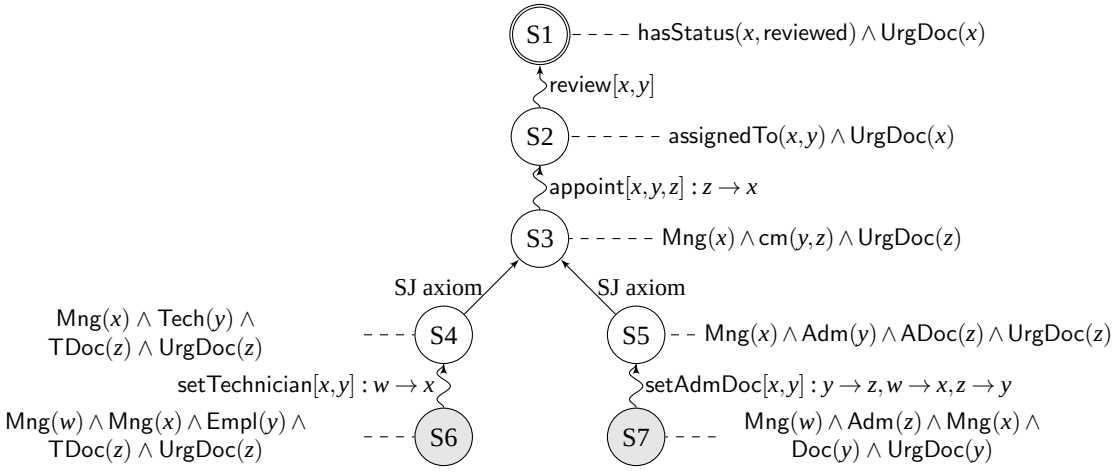

Figure 5: Abstract Planning Graph

$\{\mathsf{TechnicalDoc}(\mathsf{d001}), \mathsf{UrgentDoc}(\mathsf{d001}), \mathsf{Manager}(\mathsf{e001}), \mathsf{Administrative}(\mathsf{e003})\}$
which belongs to the set *concrete*$(S7)$ and in which we can perform the action setAdmDoc. This would lead, to the inconsistent *ABox* where d001 belongs both to TechnicalDoc and to AdministrativeDoc, which are disjoint concepts. Thus instantiated plans starting in *S7* cannot lead to the goal state.

Instances				**FP Algorithm**					**ABP+FPI Algorithm**			
Mng	Emp	TechDoc		$\lvert P_{FP}\rvert$	$\lvert V_{FP}\rvert$	*Inc*	*Time*		$\lvert P_{FPI}\rvert$	$\lvert V_{FPI}\rvert$	*Inc*	*Time*
1	1	1		3	17	13	0.06		3	7	3	0.07
1	1	2		9	38	29	0.48		5	10	4	0.30
1	1	3		25	87	66	0.28		7	13	5	0.10
1	2	2		50	154	116	0.71		10	15	4	0.15
2	2	2		80	172	134	1.35		16	16	5	0.22
2	2	3		270	413	291	3.42		22	21	6	0.18
2	3	3		816	1802	1290	33.16		33	28	6	0.24
⋮	⋮	⋮		⋮	⋮	⋮	⋮		⋮	⋮	⋮	⋮
20	20	20		-	-	-	∞		8800	862	41	197.40

Table 1: Empirical results, timing given in seconds, ∞ means more than 200 seconds.

Table 1 summarizes the experiments over different *ABox*-es, where we change the number of instances (shown in the column **Instances**). For both algorithms we indicate: (i) the size $|P|$ of the produced planning graph P (we do not consider the intermediate Abstract Planning Graph), (ii) the number $|V|$ of visited states during the creation of P (again, we do not consider the abstract states computed in the ABP phase), (iii) the number *Inc* of discarded inconsistent states, and (iv) the computation time (measured in seconds) obtained as an average of 10 runs over the same example. The timings are note very high because the code is just a prototype, but they can give an idea of the reduction of the state-space.

The results are promising because the *ABP+FPI Algorithm* performs better that the standard *Forward Algorithm*. In particular, the number $|P|$ of edges in the FP algorithm and the number of inconsistent states grows quickly with the increasing number of instances. Already 2 managers, 3 employees and 3 technical documents produce a plan with 816 edges, discarding 1,245 inconsistent states. Such a difference between the two algorithms, can be explained by the large number of redundant plans found by FP (as discussed in Examples 5 and 7).

6 Conclusions and Future Work

In this paper we have presented some preliminary work on a technique for reducing the state space in planning problems by exploiting a symbolic representation of states and reasoning techniques provided by Description Logics. Although we have chosen to adopt the DL-Lite framework and Pellet as the reasoner of our implementation, we developed the Backward State Space Reduction technique to be as independent as possible from the actual reasoning mechanism of the underlying logical representation of knowledge. The implementation of the ABP+FPI algorithm, compared to a standard Forward Planning algorithm, shows promising results both in terms of the time taken for finding the entire set of plans and in terms of the actual number of explored states.

There is a number of directions in which this work can be extended. Currently, we are working on proving the relative completeness of our ABP+FPI algorithm w.r.t. the Forward Planning algorithm, when we restrict to non-redundant plans. In the short term, we want to study the extension of actions that allows also to *remove ABox* assertions. Afterward, we plan to study the extension of our technique to the more general Datalog$^{\pm}$ family of languages as well as to allow the *creation* of new individuals as an effect of actions, thus introducing a possibly infinite planning space.

Acknowledgments. The research presented in this paper has been partially funded by the EU project ASCENS (nr.257414) and by the Italian MIUR PRIN project CINA (2010LHT4KM).

References

[1] Alessandro Artale, Diego Calvanese, Roman Kontchakov & Michael Zakharyaschev (2009): *The DL-Lite Family and Relations. J. Artif. Intell. Res. (JAIR)* 36, pp. 1–69, doi:10.1613/jair.2820.

[2] Franz Baader, Diego Calvanese, Deborah L. McGuinness, Daniele Nardi & Peter F. Patel-Schneider, editors (2003): *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.

[3] Andrea Calì, Georg Gottlob & Thomas Lukasiewicz (2012): *A general Datalog-based framework for tractable query answering over ontologies. J. Web Sem.* 14, pp. 57–83, doi:10.1016/j.websem.2012.03.001.

[4] Andrea Calì, Georg Gottlob & Andreas Pieris (2010): *Query Answering under Non-guarded Rules in Datalog+/-*. In Pascal Hitzler & Thomas Lukasiewicz, editors: *RR, Lecture Notes in Computer Science* 6333, Springer, pp. 1–17, doi:10.1007/978-3-642-15918-3_1.

[5] Diego Calvanese, Giuseppe Giacomo, Domenico Lembo, Maurizio Lenzerini, Antonella Poggi, Mariano Rodriguez-Muro & Riccardo Rosati (2009): *Ontologies and Databases: The DL-Lite Approach* 5689, pp. 255–356. doi:10.1007/978-3-642-03754-2_7.

[6] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini & Riccardo Rosati (2007): *Tractable Reasoning and Efficient Query Answering in Description Logics: The DL-Lite Family*. *J. Autom. Reasoning* 39(3), pp. 385–429, doi:10.1007/s10817-007-9078-x.

[7] Malik Ghallab, Dana Nau & Paolo Traverso (2004): *Automated Planning: Theory & Practice*.

[8] Malik Ghallab, Dana S. Nau & Paolo Traverso (2004): *Automated planning - theory and practice*. Elsevier.

[9] Babak Bagheri Hariri, Diego Calvanese, Marco Montali, Giuseppe De Giacomo, Riccardo De Masellis & Paolo Felli (2013): *Description Logic Knowledge and Action Bases*. *J. Artif. Intell. Res. (JAIR)* 46, pp. 651–686, doi:10.1613/jair.3826.

[10] John Wylie Lloyd (1993): *Foundations of Logic Programming*, 2nd edition. Springer-Verlag New York, Inc., NJ, USA.

[11] Marco Montali (2010): *Specification and Verification of Declarative Open Interaction Models. Lecture Notes in Business Information Processing* 56, Springer Berlin Heidelberg, doi:10.1007/978-3-642-14538-4.

[12] Maja Pesic, Helen Schonenberg & Wil M. P. van der Aalst (2010): *Declarative Workflow*. In Arthur H. M. ter Hofstede, Wil M. P. van der Aalst, Michael Adams & Nick Russell, editors: *Modern Business Process Automation*, Springer, pp. 175–201, doi:10.1007/978-3-642-03121-2_6. Available at `http://www.yawlbook.com/home/`.

[13] R. Reiter (2001): *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press. Available at `http://books.google.it/books?id=exa4f6BOZdYC`.

[14] Stuart J. Russell & Peter Norvig (2010): *Artificial intelligence: a modern approach (3rd ed.)*. Prentice Hall series in artificial intelligence, Prentice Hall.

[15] Evren Sirin (2006): *Combining Description Logic Reasoning with Ai Planning for Composition of Web Services*. Ph.D. thesis, College Park, MD, USA. AAI3241437.

[16] Evren Sirin, Bijan Parsia, Bernardo Cuenca Grau, Aditya Kalyanpur & Yarden Katz (2007): *Pellet: A Practical OWL-DL Reasoner*. *Web Semant.* 5(2), pp. 51–53, doi:10.1016/j.websem.2007.03.004.

Appendix

The complete specification of the Case Study.
The *TBox* is the following:

$\mathsf{Document} \sqsubseteq \neg\mathsf{Employee}$
$\mathsf{Document} \sqsubseteq \neg\mathsf{DocumentState}$
$\mathsf{DocumentState} \sqsubseteq \neg\mathsf{Employee}$
$\mathsf{Technician} \sqsubseteq \mathsf{Employee}$
$\mathsf{Administrative} \sqsubseteq \mathsf{Employee}$
$\mathsf{Manager} \sqsubseteq \mathsf{Employee}$
$\mathsf{Technician} \sqsubseteq \neg\mathsf{Administrative}$
$\mathsf{Technician} \sqsubseteq \neg\mathsf{Manager}$
$\mathsf{Administrative} \sqsubseteq \neg\mathsf{Manager}$
$\mathsf{TechnicalDoc} \sqsubseteq \mathsf{Document}$
$\mathsf{AdministrativeDoc} \sqsubseteq \mathsf{Document}$
$\mathsf{UrgentDoc} \sqsubseteq \mathsf{Document}$
$\mathsf{TechnicalDoc} \sqsubseteq \neg\mathsf{AdministrativeDoc}$
$\mathsf{Technician} \sqsubseteq \exists\mathsf{canManage}.\mathsf{TechnicalDoc}$
$\mathsf{Administrative} \sqsubseteq \exists\mathsf{canManage}.\mathsf{AdministrativeDoc}$
$\mathsf{Document} \sqsubseteq \exists\mathsf{canManage}^{-}$
$\exists\mathsf{canManage}^{-} \sqsubseteq \mathsf{Document}$
$\exists\mathsf{assignedTo} \sqsubseteq \mathsf{Document}$
$\exists\mathsf{assignedTo}^{-} \sqsubseteq \mathsf{Employee}$
$\mathsf{funct\ assignedTo}$
$\exists\mathsf{hasStatus} \sqsubseteq \mathsf{Document}$
$\exists\mathsf{hasStatus}^{-} \sqsubseteq \mathsf{DocumentState}$

The SJ axioms are:

$\mathsf{Technician}(x) \wedge \mathsf{TechnicalDoc}(y) \rightarrow \mathsf{canManage}(x,y)$
$\mathsf{Administrative}(x) \wedge \mathsf{AdministrativeDoc}(y) \rightarrow \mathsf{canManage}(x,y)$

The *ABox* is the following:

$\mathsf{Manager}(\mathsf{e001})$ $\mathsf{Technician}(\mathsf{e002})$ $\mathsf{Administrative}(\mathsf{e003})$ $\mathsf{TechnicalDoc}(\mathsf{d001})$ $\mathsf{UrgentDoc}(\mathsf{d001})$ $\mathsf{DocumentState}(\mathsf{reviewed})$

The available set Γ of actions is:

$\mathsf{appoint}[x,y,z]$: $\mathsf{Manager}(x) \wedge \mathsf{canManage}(y,z) \rightsquigarrow \mathsf{assignedTo}(z,y)$
$\mathsf{review}[x,y]$: $\mathsf{assignedTo}(x,y) \rightsquigarrow \mathsf{hasStatus}(x,\mathsf{reviewed})$
$\mathsf{setAdmDoc}[x,y]$: $\mathsf{Manager}(x) \wedge \mathsf{Document}(y) \rightsquigarrow \mathsf{AdministrativeDoc}(y)$
$\mathsf{setTechnician}[x,y]$: $\mathsf{Manager}(x) \wedge \mathsf{Employee}(y) \rightsquigarrow \mathsf{Technician}(y)$

The goal is:

$goal$: $\mathsf{hasStatus}(x,\mathsf{reviewed}) \wedge \mathsf{UrgentDoc}(x)$